exponential-e
APPLIED INNOVATION

DISASTER RECOVERY AS A SERVICE (DRAAS)

RM1557.15 G-Cloud 15 (Lot 2a Infrastructure Software as a Service)

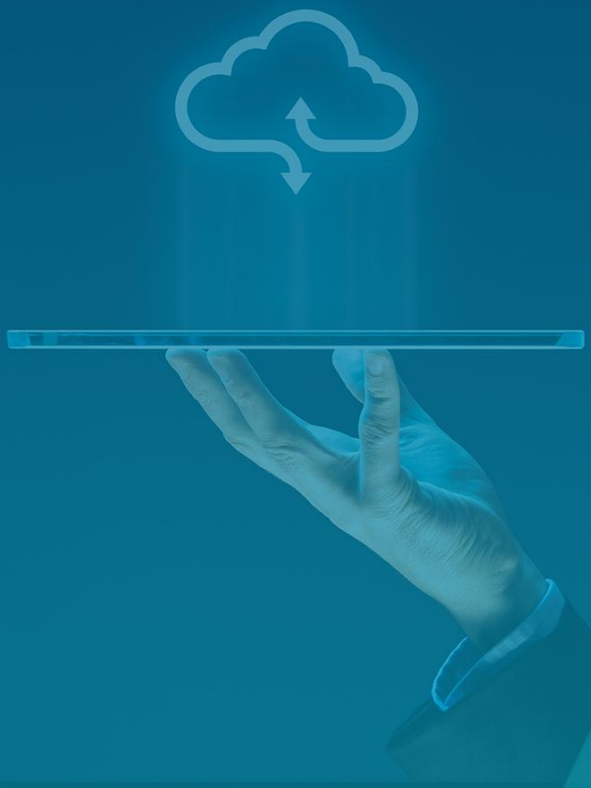

HM Government G-Cloud

Crown Commercial Service *Supplier*

bsi. ISO 9001 Quality Management | ISO 27001 Information Security Management | ISO/IEC 20000-1 Information Technology Service Management | CSA STAR Cloud Security | ISO 22301 Business Continuity Management | ISO 50001 Energy Management | ISO 14001 Environmental Management | BS 10012 Data Protection | ISO 27017 Security Controls for Cloud Services

CCS Approved Supplier:

RM1557.15 G-Cloud 15, RM3764.3 Cyber Security Services 3 DPS, RM6116 Network Services 3, RM6190 Technology Services 4, RM1043.9 Digital Outcomes (DOS) v7, RM3825 HSCN DPS, RM6094 SPARK DPS, RM6200 Artificial Intelligence DPS

PCI | CYBER ESSENTIALS CERTIFIED PLUS | HSCN

Table of Contents

Overview of the Service....2
Virtual Server Replication....3
Virtual Server Migration....3
Who Is This Service Aimed At?....3
Key Product Features....3
Service Options....4
Service Levels....4
Invocation Start Time (IST)....4
Recovery Point Objective (RPO) and Recovery Time Objective (RTO)....5
Service Credits....5
Service Management Details....5
Invocation....5
Testing....6
What Billing Models are Supported?....6
Fixed Billing....6
Pay as you Go Billing (PAYG)....6
Trial Service....6

Overview of the Service

Exponential-e™ Disaster Recovery as a Service (Server Replication) is a fully managed DRaaS service providing geo-replication of protected virtual server workloads in a number of configurations:

	Destination		
Source Environment	Microsoft Azure	VMware IaaS	Private Cloud
On-premises Hyper-V	✓	-	✓
On-premises VMware	✓	✓	✓
Exponential-e Hosted Private Cloud	✓	✓	✓
Exponential-e IaaS	✓	✓	✓
Microsoft Azure	✓	-	-
AWS EC2	✓	-	-

Our DRaaS for VMware environments is powered by Zerto, an industry-leading near continuous replication solution that integrates closely with VMware at the hypervisor layer. This provides reliable, efficient, secure and cost-effective protection for VMware workloads with RPOs1 of as little as five minutes and RTOs2 from 30 minutes.

For Hyper-V and public cloud environments, we leverage Microsoft tools including System Centre Data Protection Manager and Azure Site Recovery Manager to provide either hypervisor-layer or in-guest replication techniques, enabling us to protect virtual machines regardless of available underlying hypervisor access.

On-net customers replicating between private or IaaS environments can also benefit from seamless workload transition, without the need to change protected server IP addresses during the recovery process, made possible by our investment in Software Defined Networking (SDN). These techniques can also eliminate risk from permanent workload migration between platforms.

[1] The amount of time reflecting the amount of data that potentially could be lost during a disaster recovery.
2 The targeted period of time within which the service will be restored.

What Can I Do with This Service?

Virtual Server Replication

Virtual servers can be seamlessly replicated to a second data centre in order to provide recoverability in the event of loss of, or loss of access to, the primary data centre.

Virtual Server Migration

Virtual servers protected with our DRaaS service can be permanently moved to the target environment in order to decommission the primary environment, for example as server hardware becomes end-of-life. In this scenario, DRaaS protection can then be re-established to another secondary data centre.

Who Is This Service Aimed At?

- ✓ Customers with virtualised environment requiring DRaaS with an op-ex pricing model.
- ✓ Customers requiring DRaaS with cross hypervisor replication capability.

Key Product Features

Feature	Benefit
Secure & Sovereign	UK Sovereign service options, run from tier-3 data centres with ISO 27001 certification. Data encrypted options both in-transit and at rest.
Predictable Billing	Per-Protected-VM and per-GB Billing No data transit charges No API charges
Scale Indefinitely and On-Demand	Eliminates the need for fork-lift storage expansions
Multi-Hypervisor Support	A single DRaaS service that can cover multiple environments in multiple locations
Storage Agnostic	No storage vendor dependencies, licensing, or tie-in

Application Agnostic	Replication at the hypervisor layer enables the protection of all applications and guest operating systems
Low RTPOs	Zerto based replication for VMware environments can achieve near continuous replication without the need for VMware snapshots, providing RPOs as short as five minutes As a managed service, we offer an SLA over the time to invoke recovery of 30 minutes.
Defined Failback	Automatic failback processes Pre- and Post- failover scripting for recovery customisation
Application Groups	Consistent recovery for multi-tier applications VM group boot ordering and automatic re-IP process
Isolated Testing	No-impact failover testing to isolated environments decouples the testing of critical infrastructure in the failover path.

Service Options

The service deployment specifics are chosen dependent on the source and destination platforms in-scope, and the available hypervisor access.

For environments replicating to the Exponential-e IaaS platform, customers are able to seed virtual machine images in order to reduce network bandwidth requirements during the set-up phase.

Service Levels

We will use reasonable endeavours to ensure that the availability of the service purchased by the customer in a given calendar month equals the applicable Availability Commitment.

To define availability, we monitor a number of service elements — some generic, some service specific – which collectively enable the customer to use or access the service. If the availability of the service is less than the associated Availability Commitment, the customer may request Service Credits for the service within 30 calendar days of the service being deemed unavailable. Service credits are measured in the table below, for more information on our Service Level Agreements (SLAs), please see the Exponential-e SLA definition document.

Invocation Start Time (IST)

The DRaaS service is subject to the following target IST:

	IST*
Server Replication Failover Invocation Time	30 Minutes

** The time taken for Exponential-e to start invocation of the Clone Server in accordance with the Invocation Initialization Document measured from the raising of the trouble-ticket.*

Recovery Point Objective (RPO) and Recovery Time Objective (RTO)

RPO and RTO (combined, RTPO) will depend upon a number of customer-specific factors and therefore cannot be detailed within this document.

Service Credits

Metric	Measure	Service Credit*
Invocation Start Time	> 1 hour	5%
	> 1 hour 30 minutes	10%
	> 2 hours	15%
	> 2 hour 30 minutes	20%

* The service credit is applied as a percentage of the Monthly Charge for the Service for the month concerned (whether based on the fixed Annual Charge, Pay As You Go Charges or, in the case of a Hybrid Billing Model, both).

Service Management Details

Invocation

The customer's authorised personnel must contact us and request invocation of the DRaaS service. We will carry out the invocation plan as documented in the Invocation Initialisation Document. Additional resource charges will arise for metered platform resources, as detailed in the separate Pricing Document, will be incurred for the duration of the invocation.

We have multiple customers using our DRaaS service, using the same underlying resources. Accordingly, we cannot guarantee that there will not be competing demands in the event of multiple simultaneous disasters or a major disaster affecting a number of customers. If such disaster(s) occur, all invocation requests will be enacted on a 'first come, first served' basis. We shall not be liable if, due to such disaster(s), invocation of the service cannot be commenced within the Invocation Start Time set out above.

Testing

The customer will receive a test allowance of two days per calendar year, or such greater amount as is stated in the Agreement, for testing the DRaaS service.

Testing will be scheduled and conducted in accordance with our policies and procedures then in effect, as amended from time to time. The test environment provided will be an isolated environment and access to the virtual machines shall be provided either by portal console access or with a bastion host, as defined in during the initial delivery.

Full failover tests can also be provided, subject to additional Professional Services costs as scoped as part of the service discovery exercise.

We may cancel a test in order to undertake a live invocation of another customer's DRaaS service. We will use reasonable endeavours to reschedule cancelled tests, but no allowances or credits will be given.

What Billing Models are Supported?

Fixed Billing

The customer is charged a fixed Annual Charge and provided access to a fixed resource pool. Fixed elements typically include Replication Licenses and Replicated Server Storage.

Pay as you Go Billing (PAYG)

The customer is charged for the actual resources used in accordance with the Rate Card and the applicable Service Definition and will be billed monthly in arrears. Actual usage levels will be metered and recorded by Exponential-e.

PAYG elements include IaaS platform metered resources, such as vCPU, vRAM, Disk, or Azure Virtual Machines used for failover testing or during a live invocation of service. PAYG can also be used for additional ad-hoc resource consumption, such as additional replicated server storage.

Trial Service

Due to the integration and work required no trial is available for our DRaaS solutions. If a trial is required, we are able to provide a one month paid Proof of Concept (POC). On completion of the POC, customers choosing to continue with the service will be refunded the one-month trial period.

